arXiv:1704.07863v1 [cs.CV] 25 Apr 2017

Multi-View Dynamic Facial Action Unit Detection

Andrés Romero Juan León Pablo Arbeláez

Universidad de los Andes

Abstract

We propose a novel convolutional neural network architecture to address the fine-grained recognition problem of multi-view dynamic facial action unit detection. We leverage recent gains in large-scale object recognition by formulating the task of predicting the presence or absence of a specific action unit in a still image of a human face as holistic classification. We then explore the design space of our approach by considering both shared and independent representations for separate action units, and also different CNN architectures for combining color and motion information. We then move to the novel setup of the FERA 2017 Challenge, in which we propose a multi-view extension of our approach that operates by first predicting the viewpoint from which the video was taken, and then evaluating an ensemble of action unit detectors that were trained for that specific viewpoint. Our approach is holistic, efficient, and modular, since new action units can be easily included in the overall system. Our approach significantly outperforms the baseline of the FERA 2017 Challenge, which was the previous state-of-the-art in multi-view dynamic action unit detection, with an absolute improvement of ***14%****.*

1. Introduction

The field of human facial expression interpretation has benefited from seminal contributions by renowned psychologists such as P. Ekman, who characterized and studied the manifestation of prototypical emotions through changes in facial features [11]. From the computer vision perspective, the problem of automated facial expression analysis is a cornerstone towards high-level human computer interaction, and its study has a long tradition within the community. Initially, the problem was approached by focusing on its most basic version, and classifying static images or short sequences of faces into a handful of prototypical emotions (*e.g.*, happiness, sadness, fear, etc.). Recent methods[24, 17] have achieved significant progress in the study of this task, leading to the near saturation of standard benchmarks such as the CK+ dataset [34].

In order to systematize the study of facial expressions,

Figure 1: **Results of our approach.** Given a video of a human face that was taken from an unknown viewpoint, our system predicts the presence or absence of multiple action units in each frame.

Ekman and his collaborators designed the Facial Action Coding System (FACS) [12]. FACS relies on identifying visible local appearance variations in the human face, called Action Units (AUs), that are produced by the individual activations of facial muscles. Thus, a raised eyebrow is coded as the activation of the outer brow raiser, and noted AU02. Since any facial expression can be represented as a combination of action units, they constitute a natural physiological basis for face analysis. The existence of such a basis is a rare boon for a computer vision domain, as it allows focusing on the essential atoms of the problem and, by virtue of their exponentially large possible combinations, opens the door for studying a wide range of applications beyond prototypical emotion classification. Consequently, in the last years, the main focus of the community has shifted towards the detection of action units, and recent datasets such as BP4D [51] and the FERA 2017 challenge [43] include annotations by trained psychologists for multiple action units in individual video frames.

When compared to global emotion classification, action unit detection is a much challenging and fine-grained recognition task, as shown by the local and delicate appearance

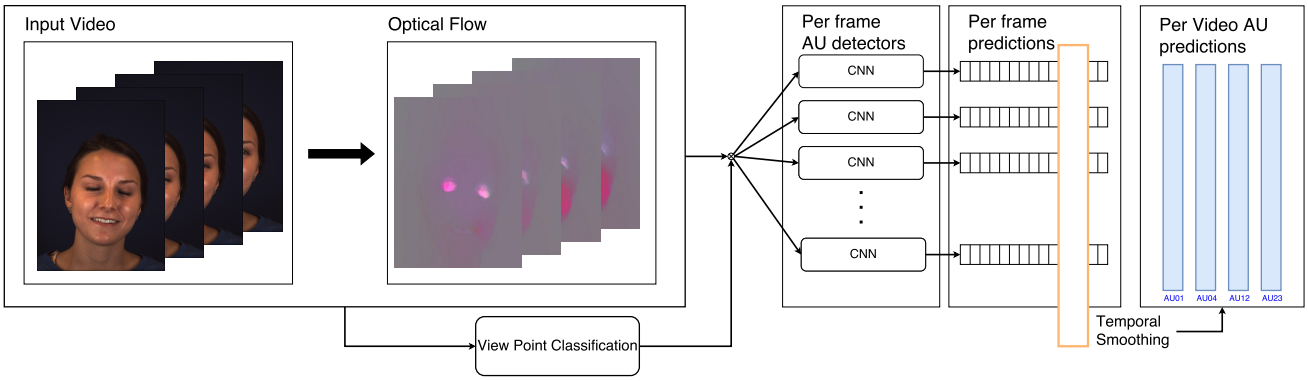

Figure 2: **Overview of AUNets.** Our system takes as input a video of a human head and computes its optical flow field. It predicts the viewpoint from which the video was taken, and uses this information to select and evaluate an ensemble of holistic action unit detectors that were trained for that specific view. Final AUNets predictions are then temporally smoothed.

variations in Fig. 1. The expression of action units is typically brief and unconscious, and their detection requires analyzing subtle appearance changes in the human face. Furthermore, action units do not appear in isolation, but as elemental units of facial expressions, and hence some AUs co-occur frequently while others are mutually exclusive. In response to these challenges, the literature has converged to a dominant approach for the study of action unit detection [54, 53, 50, 21, 13, 1, 25, 16, 32] that uses the localization of facial landmarks as starting point. Recent methods implement this paradigm in two different ways, either using facial landmarks as anchors and analyzing the appearance of patches centered in those keypoints as well as their combinations, or using landmarks to perform face alignment prior to AU detection. Focusing on fixed facial regions has the advantage of constraining appearance variations. However, facial landmark localization in the wild is still an open problem, and its study requires also very expensive annotations. Furthermore, while such an approach is suitable for a near frontal face setup, as is the case in the BP4D dataset, it finds its limitations in a multi-view setting such as FERA 2017, in which large variations in head pose imply occlusion of multiple facial landmarks and significant appearance changes.

In this paper, we depart from the mainstream approach for AU detection and propose a system that directly analyzes the information on the whole human face in order to predict the presence of specific action units, bypassing thus landmark localization on the standard benchmarks for this task. For this purpose, we leverage recent insights on designing Convolutional Neural Networks (CNNs) for recognition applications. First, we observe that large-capacity CNN architectures [26, 40, 41, 18] that were originally designed for object recognition on datasets such as Imagenet [10] analyze a whole raw image and are capable of making subtle category distinctions (*e.g.* between dog breeds). Moreover, they can be successfully specialized for other fine-grained recognition problems [47, 5]. Therefore, we formulate the problem of predicting the presence or absence of an specific AU in a single face image as holistic binary classification. We explore the design space of our approach and, in particular, the trade-off between efficiency and accuracy when considering shared or independent convolutional representations for different action units.

Second, since action units appear dynamically within the spatio-temporal continuum of human facial expressions, we model explicitly temporal information in two different ways. At frame level, we compute an optical flow field and use it as an additional cue for improving AU detection.Then, we enforce short-term consistency across the sequence by a smoothing operator that improves our predictions over a small temporal window.

Third, in order to address multi-view detection, the main focus of the recent FERA 2017 challenge, we take inspiration from the work of Dai *et al.* [18], who showed that training CNNs for a hard recognition task such as object instance segmentation on MS-COCO [28] benefits from decomposing learning into a cascade of different sub-tasks of increasing complexity. Thus, our multi-view system starts by predicting the overall view of an input sequence before proceeding to the finer grained task of action unit detection. Furthermore, as the experiments will show, our system benefits from a gradual domain adaptation to the final multi-view setup. Fig. 2 presents an overview of our approach, which we call AUNets.

We perform an extensive empirical validation of our approach. We develop our system on the BP4D dataset, the largest and most varied available benchmark for frontal action unit detection, where we report an absolute improvement of **8%** over the previous state-of-the-art by Li *et*

al. [27]. We then turn to the FERA2017 challenge to evaluate our multi-view system, and report an absolute improvement of **14%** over the challenge baseline of [43]. In order to ensure reproducibility of our results and to promote future research on action unit detection, all the resources of this project –source code, benchmarks and results– will be made publicly available.

2. Related Work

Before deep learning techniques became mainstream within the field of computer vision [26], most methods relied on the classical two-stage approach of designing fixed handcrafted features such as SIFT [33] or LBP [37], and then training unrelated classifiers for recognition [44, 31, 1, 50, 13, 53, 25]. However, similarly to AUNets, the best performing techniques currently available [21, 16, 54, 27] rely on the power of deep convolutional neural networks for joint representation and classification.

Most recent methods [30, 1, 50, 13, 53, 21, 16, 54, 25] follow the paradigm of first detecting facial landmarks using external approaches such as Active Appearance Models [8], either to treat these keypoints as anchors for extracting rectangular regions for further analysis, to perform face alignment, or both. The recent method of Li *et al*. [27] does not require robust facial keypoint alignment as it is trained taking this issue into account; however, facial alignment is recommended and the method is developed and tested only in a frontal face setup. In contrast, AUNets operate on the whole face and do not require any particular alignment in existing AU detection multi-view benchmarks.

Pioneering methods for this task used the whole face image as input [44, 29, 31]. However, the trend reversed towards analyzing patches in more local approaches such as [30, 1, 50, 13, 53, 21]. State-of-the-art techniques [16, 54, 27] join AUNets in returning to a holistic face analysis, in particular Li *et al*. [27] shares ideas with Zhao *et al*. [54] in forcing a CNN-based approach to specifically focusing on specific regions of the face by using a map saliency. Our method is far from these approaches since we train specific networks for each AU which avoids the need of building attention maps.

Given that groups of action units can co-occur or be mutually exclusive in the human face, several methods [54, 44, 53, 27] have approached the task as multi-label learning. However, as AUs databases are becoming larger and better annotated [52, 35], these methods need to be completely retrained to integrate new action units. In contrast, AUNets are modular by design, and can naturally evolve towards more general datasets and new action units. Furthermore, by analyzing the whole face, AUNets learn naturally the relevant local face regions for predicting each action unit.

While most methods operate on single RGB images both at training and testing[31, 1, 16, 29], some techniques focus on exploiting the temporal information for improved AU detection[50, 25, 30]. In particular, Jaiswal *et al*. [21] and Chu *et al*. [6] use CNN's and Bidirectional Long Short-Term Memory to model time dependencies. Similarly, Liu *et al*. [32] tackle the temporal information by taking advantage of the Optical Flow [19]; for this purpose, starting from the OF they extract hand-crafted features as histogram oriented [4] in order to train a classifier. However, our method takes advantage of the full resolution of the OF by feeding it into a CNN. Moreover, on behalf of the Optical Flow, we perform simple statistical operations to smooth the temporal flow as a post-processing step.

Since the problem of Facial Expression recognition is evolving throughout new dataset, FERA17 Challenge [43] change the way how it approach was perceived. That is not only to analyzing frontal images but multi-view rendered facial images [43]. Following this change in paradigm, AUNets can easily morph to this novel paradigm by using an intermediate view detector and then more specialized AU detectors regarding the view, instead of automatically frontalizing each facial view [43] which is carried out by facial alignment.

3. Approach

3.1. Facial Expression Representation

We formulate individual action unit detection as holistic binary classification in order to build on recent high-capacity CNN architectures [26, 40, 41, 18] that were originally designed for object classification in large-scale datasets such as Imagenet [10]. Recent works have shown that these networks can learn useful representations for different object-level applications such as contour detection [45, 46], semantic segmentation [36], instance segmentation [9], and action recognition [3].

However, relative to ImageNet, the domain adaptation we require for action unit detection is much larger than the one for the above mentioned problems, which are typically studied in VOC PASCAL [14]. Therefore, we start by learning a high-level representation on a related but simpler facial expression analysis task. We achieve this goal by converting the fully connected layers of an ImageNet architecture such as VGG [40] into convolutional layers, and then re-train it for the traditional task of prototypical emotion classification. In order to generalize across scale, ethnicity, gender and age, we train the encoder jointly on several standard emotion recognition datasets with different characteristics. After learning is completed, we remove the output layer of the network and use the learned weights as convolutional encoder for facial expressions in all subsequent experiments. Our results show that this initial domain adaptation is critical for good performance, indicating that the convolutional encoder is indeed learning features that are

specific to faces.

3.2. Single-Image Action Unit Detection

Starting from the learned convolutional encoder, training a specialized detector for a given action unit in single RGB images is straightforward. We extend the encoder with a binary output softmax layer predicting the presence or absence of the AU in the image, we initialize it with random weights, and then resume training in a database that provides AU annotations, such as BP4D. However, when considering multiple action units, two architectural options emerge, depending on whether or not we share features at the lower layers of the network.

Hydra-Net The first option is to preserve a shared representation and to train individual modules for different AUs. This strategy is implemented by freezing the lower layers of the convolutional encoder and learning separate weights at the upper layers sequentially for each action unit. We call this network "Hydra-Net" in reference to the fabulous Lernaean Hydra from the Greek mythology, a dragon with a single body and multiple heads that would multiply when severed. Hydra-Net is efficient at test time and modular by design, as it can be easily extended to new action units by "growing" additional heads. Furthermore, as the experiments will show, this baseline system already outperforms the state-of-the-art on BP4D for single-image AU detection.

AU-Nets The elegance of a shared representation in Hydra-Net comes at the price of a reduction in absolute performance, because only a fraction of the original encoder weights are being specialized for each action unit. In order to quantify this trade-off, we train a second variant of our system in which we relearn all the weights of the encoder for each action unit. This strategy results in a battery of independent AU detectors, which we call "AU-Nets". When compared to Hydra-Net, the larger-capacity ensemble of AU-Nets requires more data at training time and is less efficient at test time. However, both approaches are modular in the presence of new action units and, as the experiments will show, the increased flexibility of AU-Nets allows them to focus better on specific regions of the face associated to different action units, leveraging thus local information for improved performance.

3.3. Dynamic Action Unit Detection

Although trained human observers can annotate action units on single images, AUs are continuous and precisely localized events that occur within the spatio-temporal flow of facial expressions, and therefore an explicit analysis of motion information should facilitate AU detection when video data is available. We model motion by computing a dense optical flow field and embedding it into a three dimensional space in which the first two dimensions are normalized x and y pixel motion and the third dimension is the optical flow magnitude. This embedding is common in video analysis applications [15, 38] and has the advantage of providing a motion signal that has the same scale and dimensions as the original RGB frame. We considered three architectures for combining color and motion information, as illustrated in Fig. 3.

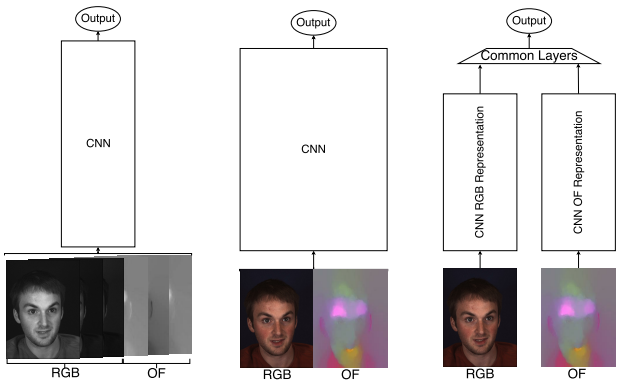

Figure 3: We consider three architectures for combining color and motion information: **(a)** Additional input channels, **(b)** an extended image, and **(c)** two separate CNN streams. See text for more details.

Additional input Channels. A first option is to consider the optical flow embedding as three additional channels of the input RGB image. This architecture is economical, as it adds only additional weights for the first layer of the network, which are initialized as copies of the RGB weights. This first layer is then learning to combine all channels for each spatial location.

Concatenation. A larger architecture is obtained by concatenating color and motion images in one of the two spatial dimensions. In this case, the network analyzes the two signals with the same first-layer filters, but its fully convolutional architecture is expanded accordingly in order to combine their information; resulting also in larger capacity models.

Two streams. A third option is to consider two different streams to represent color and motion information, and to merge them in higher layers. In this case, we freeze the weights of the RGB encoder and copy them to initialize the motion encoder. We are thus adapting the RGB representation to the motion signal, while simultaneously learning its optimal combination with the color representation in the

deeper layers. This strategy is also known as 'π Network'

The experimental section will show that explicit modeling of motion between successive frames is beneficial for AU detection. Nevertheless, we note that our dynamic AU detectors above still operate independently in each frame, and hence cannot capture the temporal smoothness of action unit activation. In order to enforce such consistency, we perform a simple post-processing by applying a temporal sliding median filter over predictions of each action unit across the video. We found this simple smoothing to be empirically beneficial for all the variants of our approach.

3.4. Multi-View System

Our final system is evaluated on the FERA 2017 challenge, which focuses on multi-view AU detection. Starting from a video of a face that was taken from an unknown viewpoint among nine options, the task consists in predicting labels for twelve common action units on all its frames. We extend our system in two ways to address this multi-view version of the problem. First, we train a view classifier, which takes the video as input and predicts its viewpoint. Once the viewpoint has been estimated, our cascaded approach proceeds to evaluate an ensemble of temporally consistent dynamic AU-Nets that were trained for that specific view. In order to perform domain adaptation, we first retrain our system developed in BP4D in the frontal view of FERA 2017, and then use those weights to initialize learning for the other views.

4. Experimental Validation

In this section, we provide the empirical evidence to support our approach. We first conduct extensive experiments in the BP4D dataset analyzing different aspects of our method for frontal faces. We then turn to the FERA 2017 dataset in order to extend our system to a challenging multi-view setting.

4.1. Results on BP4D-Spontaneous

Experimental Setup

The BP4D-Spontaneous database [51] is currently one of the largest and most challenging benchmarks available for the task of action unit detection [51]. BP4D contains 2D and 3D videos of spontaneous facial expressions in young people. All of them were recorded in non-acted scenarios, as the emotions were always elicited by experts. This dataset contains around 150.000 individual frames from 41 subjects annotated by trained psychologists.

The BP4D dataset has a significant class imbalance for most AUs; for instance, there are 6 times more negatives samples than positive samples for AU23, while AU12 is almost balanced. We follow the common practice in the literature [54, 53, 6, 27], and consider the 12 most frequent AUs [51] for the BP4D dataset. In order to compare our approach directly with the state-of-the-art, we report results using the standard performance metric from the literature, the image-based F1-score (f1-frame) [42], also known as F-measure and defined by the harmonic mean of precision and recall, which considers the optimal operating regime of the detector. For all the BP4D experiments, we strictly follow the experimental framework of [54] and [6] and perform three-fold cross validation using their same splits, in order to ensure that our results are directly comparable.

Domain Adaptation

We first quantify the importance of learning a high-level representation for faces prior to action unit detection. We select three popular CNN architectures that were originally designed for ImageNet and that have been applied to multiple vision problems: AlexNet [26], VGG-16 [40], and GoogLeNet [41]. In order to expose the network to a large variety of faces, we train on the union of three standard datasets for emotion classification with very different characteristics: CK+ [34], RafD [23] and Bosphorus [39]. We modify the output layer for 8-way emotion classification and then train the network for 5 epochs with a learning rate of 10^{-4} and weight decay of 0.005.

We then use these six convolutional encoders as different initializations to the Hydra-Net architecture in BP4D and train them until convergence. To overcome the dataset's AU skew, we augment the positives training samples using the well known strategy of image jittering [26], until we reach an approximate balance with the negative samples.

Results are presented in the left panel of Table 1. We first observe that all three architectures produce reasonable AU detection results when initialized with ImageNet weights, and that deeper architectures produce improved results, in accordance to their original performance in ImageNet. This behavior contrasts with random initialization, which causes the training process to diverge for the three CNNs (results not included). Furthermore, we observe a significant improvement in AU detection for all base CNN architectures when initializing with our face expression convolutional encoder instead of the original ImgeNet weights. Among the three CNNs considered, VGG-16 generalizes better to our application domain, obtaining a $+15\%$ absolute improvement in performance over ImageNet weights. These results highlight the relevance of domain adaptation for facial expression analysis. We use the convolutional encoder based on the VGG-16 architecture for all subsequent experiments.

Shared vs. Independent Representation

We next explore the trade-off between efficiency and accuracy when considering shared or independent represen-

tations for different action units. For this purpose, we start from our shared representation Hydra-Net and train with the above mentioned parameters a new fully independent detector for each AU. Learning is carried out by unfreezing all layers in the convolutional encoder. The central column of Table 1 presents the results for independent detectors, which we call in the sequel AUNets. In comparison to HydraNet, under an independent representation 10 out of the 12 action units show improvements. Moreover, we observe an overall better performance under this strategy as, on average, the F1 measure increases by 3.3 over the whole set.

Optical Flow

To incorporate short term motion patterns into our experimental setup, we estimate the optical flow (OF) with the variational model proposed by Brox *et al*. [2]. Given the large scale of the BP4D and FERA17 datasets, we resort to a GPU implementation of this algorithm[20] as it provides a fair trade-off between accuracy and computation time. In practice, we resize the images to 512x512, which allows to calculate the OF for the whole training set of FERA17 in under 60 hours with a single Titan X GPU.

To adapt the resulting dense vector map into the AU-Nets, it is transformed into an RGB image as proposed by [15, 38], by independently normalizing the x and y dimension of the OF vectors, and then rescaling to 255 to conform the images R and G channels, whereas the B channel is calculated as the norm of the original vector.

Initially, we consider motion information in isolation and train our AUNets ensemble under the same experimental set up as the original RGB frames. Column 'OF' in the right panel of Table 1 summarizes the results for this experiment. Even without any color information, the motion signal alone achieves a performance of 52.8%, which is competitive with the state-of-the-art in BP4D.

We then explore the strategies outlined in Section 3.3 to fuse information from the OF and RGB domains. Table 1-right shows the results for the three proposed architectures, including two options for the Two Stream strategy. We observe that the motion signal is complementary to the color information, as the inclusion of OF always brings an improvement over the RGB results. The Two Stream and the Concatenation strategies are both superior to the Additional Channels, and provide comparable performance. We choose the latter because its number of learnt parameters is 13% smaller than the former, which means that training will be faster and overfitting will be less likely.

We conclude that, overall, the best approach is to concatenate the OF and RGB information in a vertical arrangement; it has the very same performance as a π network joint at the FC6 layer, yet, the number of learnt parameters is 13% smaller, this means that further domain adaptation will be faster, and overfit will be less likely.

Temporal Smoothing

Since action units appear in smooth, short-term duration intervals, we add a post-processing step which aims at removing instantaneous AU activations and possible misclassifications inside a continuous sequence. For this goal we use a median operator in a sliding window fashion over the AU detection the sequences.

Our approach uses small window sizes as we want to avoid the suppression of properly detected but short temporal series of AU activations. Therefore, we test with window sizes: {3,5,7,9,10,11}, these remain fixed for all AUs over the entire train set. Empirically we conclude that the best window sizes are {3,5,7} with 5 being the optimal, longer temporal windows do not improve the base performance. Table 1-right shows the result (+median) of performing this temporal smoothing to the best OF arrangement (Vertical).

Comparison against the state-of-the-art

We now compare our results against the state-of-the-art approaches presented on [27, 6, 54, 53]. Table 2 summarizes the average of three-fold cross validation for each action unit. We observe that our method consistently outperforms all methods for 10 out of 12 action units, with an average improvement over the F1 measure of **8%**, over EAC.

4.2. Results on the FERA 2017 Challenge

Experimental Setup

The new version of the FERA17 challenge [43] introduces a novel way to approach the problem of facial expression analysis. This dataset contains a 3D video database rendered mainly from BP4D [51]. Additionally, it includes subsets from a non publicly available database [52]. The FERA17 dataset includes 9 different camera angles for non-symmetric facial views. It contains videos from 41 subjects and around 1.500.000 frames for the whole training split. An additional set of videos from 20 different subjects with about 750.000 frames integrate the validation split. A final set of 1080 videos are included in the withheld test split.

Viewpoint Classification

Our AU-Net ensemble is view specific, and thus relies on a proper view selection for optimal results. We approach this problem by performing a view recognition subtask, prior to the AU-Net ensemble.

For this sub-task, we build a view classifier on top of another deep convolutional encoder. We start from the Caffe-GoogleNet[22] reference model since we require a simple

AU	AlexNet		GoogleNet		VGG		VGG	VGG+OF					
	ImgNet	EmNet	ImgNet	EmNet	ImgNet	EmNet	AUNet	OF	Channels	π/fc6	π/fc7	Vertical	+Median
1	19.8	40.3	27.0	34.1	27.4	**43.0**	51.1±2.2	35.4±6.7	50.1±1.1	51.4±2.4	53,4±5,6	52.4±1.6	**55.3±2.5**
2	16.4	23.3	17.2	28.0	27.6	**30.7**	45.7±9.2	29.8±6.4	45.6±8.5	45.6±8.2	49,8±3,4	45.8±8.7	**51.0±9.5**
4	31.4	43.3	44.9	46.0	40.0	**48.8**	52.1±4.5	38.1±0.2	51.1±0.9	52.9±4.6	**56.0±3.2**	55,5±3,1	53.9±4.9
6	49.4	65.3	71.6	76.3	57.2	**77.2**	77.5±1.9	70.9±5.1	77.6±1.4	77.9±1.6	77.4±1.9	77,4±1,5	**78.2±1.7**
7	51.3	71.3	64.8	70.5	62.2	**76.6**	77.6±4.4	70.7±3.7	78.3±5.6	78.1±4.5	78.4±3.0	78,2±3,1	**78.4±4.7**
10	47.2	83.7	70.4	75.0	47.1	**78.0**	83.0±2.4	76.6±2.9	83.3±2.3	84.0±1.7	83.2±2.3	83,0±2,3	**84.2±1.7**
12	53.7	79.5	75.1	84.1	58.8	**87.2**	86.8±2.3	77.3±1.7	86.9±2.3	86.9±2.3	**88.0±0.2**	87,9±0,5	87.3±2.2
14	23.4	49.1	51.1	58.4	59.9	**68.2**	67.4±4.8	64.0±6.0	67.6±4.6	**68.0±4.4**	67.3±4.5	66,8±3,8	67.3±5.8
15	22.7	40.0	31.3	32.3	28.4	**41.1**	42.0±3.4	34.9±4.5	44.9±4.9	46.5±5.1	44.2±5.4	46,4±9,0	**47.6±4.4**
17	21.4	27.3	51.2	54.9	41.2	**61.0**	61.8±3.9	56.5±2.6	62.7±3.7	62.8±4.5	61.5±3.0	61,1±3,1	**63.9±4.4**
23	20.4	34.2	28.1	29.6	32.5	**35.1**	40.9±2.9	36.7±2.1	43.5±3.0	43.4±3.7	41.7±3.3	41,1±3,2	**44.7±4.4**
24	21.0	43.1	37.5	39.7	29.8	**44.6**	50.9±1.5	42.1±0.5	51.7±2.7	53.4±1.7	52.7±1.1	52,4±0,2	**54.2±2.2**
Av.	31.5	49.3	47.5	52.2	42.7	**58.1**	61.4±3.6	52.8±3.5	61.9±3.4	62.7±3.7	62.3±3.2	62,7±3,2	**63.8±4.0**

Table 1: **Control experiments for convolutional encoder**. Results are reported over BP4D by using 3-fold cross validation using the F1-metric. The left panel compares several base networks as starting point for training Hydranets. The central and right panels present experiments on AUNets and different combination of Optical Flow and Color.

AU	JPML[53]	DRML[54]	MSTC [6]	EAC[27]	Ours
1	32.6	36.4	31.4	39.0	**55.3**
2	25.6	41.8	31.1	35.2	**51.0**
4	37.4	43.0	**71.4**	48.6	53.9
6	42.3	55.0	63.3	76.1	**78.2**
7	50.5	67.0	77.1	72.9	**78.4**
10	72.2	66.3	45.0	81.9	**84.2**
12	74.1	65.8	82.6	86.2	**87.3**
14	65.7	54.1	**72.9**	58.8	67.3
15	38.1	33.2	34.0	37.5	**47.6**
17	40.0	48.0	53.9	59.1	**63.9**
23	30.4	31.7	38.5	35.9	**44.7**
24	42.3	30.0	37.0	35.8	**54.2**
Av.	45.9	48.3	53.2	55.9	**63.8**

Table 2: **3-fold cross validation for each Action Unit over BP4D dataset using F1-metric.** Comparisons against the state-of-the-art methods: JPML [53], DRML[54], MSTC [6], and EAC [27].

architecture to train, then we proceed to learn a representation using the full training set, freezing the first 6 Inception modules and learning over the final 3, we also drop the two deep supervision branches as they remain connected to frozen segments of the encoder. We learn from scratch the final fully connected layer to produce a 9-dimensional output. This viewpoint-network is trained for two epochs, with an initial learning rate of 5^{-5}, gamma parameter of 0.1, and a weight decay of 5^{-5}, we make learning rate steps each epoch. The network is trained until convergence.

As final step we get a single prediction for any video as $\max_v(p)$, where p is the per frame prediction and v is the full set of frames for a video. Evaluation results in table 3 suggest that the viewpoint classifier is almost perfect in the FERA 2017 setup. This result indicates that AU detection will not be affected in any significant way by prior view classification.

Multi-View System

For our first approach to the FERA17 Challenge, we want to assess the generalization capability of our system, so we proceed to evaluate our system trained in BP4D over the entire validation partition of FERA17 without any retraining. The evaluation procedure results in an performance of **46.5%**, which shows that our system trained on BP4D with no previous training on FERA17 already surpasses the baseline approach for validation partition in FERA17 Challenge [43] that is 40.4%. We observe that, when the view is close to the frontal one, our model improves performance, which is 51.9% over the frontal view, and 41.4% over an upper-left view.

we use the BP4D models as starting-point to train over the frontal view of FERA17 dataset. We follow the same strategy for each view, in order to learn the entire 9 views for 10 AUs, resulting in a total of 90 models. For each view, we learn our models with the same setup as in BP4D, as outlined in Section 3.4. We train each view and each AU until convergence with learning rate of 10^{-4} and weight decay of 5^{-3}, the whole system is trained for a total of 50h in a single Titan X GPU. AUNets+OF trained on FERA17 and evaluated in the validation set obtain **63%** on average, which widely outperforms the 40.4% of the baseline [43].

	View 1	View 2	View 3	View 4	View 5	View 6	View 7	View 8	View 9	Average
Per Frame Precision	0.96	0.96	1.0	0.96	0.97	0.99	0.98	0.91	0.94	0.96
Per Frame Recall	0.99	0.95	0.95	1.0	0.96	0.98	0.95	0.93	0.96	0.96
Per Video Precision	0.97	0.96	1.0	0.96	0.97	0.99	0.98	0.92	0.94	0.97
Per Video Recall	0.99	0.96	0.95	1.0	0.96	0.98	0.95	0.94	0.96	0.97

Table 3: **Evaluation of Multi-view Classifier**, Per frame and per video results for the view recognition subtask in the validation set of FERA17.

AU	F1-score (Baseline [43]) Ours									
	v1	v2	v3	v4	v5	v6	v7	v8	v9	**Global**
1	(10.3) **45.1**	(15.0) **47.3**	(13.6) **45.2**	(19.3) **54.5**	(19.6) **49.5**	(17.1) **54.8**	(18.0) **45.0**	(14.5) **49.6**	(12.3) **47.1**	(15.5) **48.7**
4	(15.0) **53.4**	(15.9) **57.7**	(14.8) **46.0**	(19.1) **53.3**	(19.0) **81.5**	(18.3) **55.3**	(20.2) **59.1**	(20.2) **51.1**	(17.5) **43.1**	(17.8) **55.6**
6	(50.5) **75.4**	(55.7) **73.6**	(13.4) **77.6**	(68.9) **77.2**	(74.7) **77.7**	(72.4) **79.0**	(56.0) **77.3**	(53.2) **72.3**	(49.3) **74.4**	(54.9) **76.1**
7	(72.1) **81.6**	(72.9) **81.9**	(41.3) **81.4**	(74.6) **81.9**	(79.7) **83.3**	(78.7) **82.6**	(71.6) **80.7**	(74.7) **81.7**	(75.8) **80.8**	(71.3) **81.8**
10	(55.4) **82.4**	(71.0) **83.9**	(64.2) **82.6**	(77.7) **83.1**	(77.6) **84.5**	(75.0) **84.8**	(63.9) **85.0**	(67.9) **83.7**	(65.9) **82.0**	(68.7) **83.6**
12	(52.2) **86.2**	(67.8) **86.6**	(18.4) **84.6**	(78.6) **86.7**	(80.9) **85.3**	(77.1) **85.2**	(59.6) **85.2**	(63.8) **84.9**	(60.1) **82.9**	(62.1) **85.3**
14	(51.5) **72.3**	(56.3) **73.6**	(9.0) **71.7**	(67.5) **73.8**	(72.4) **77.0**	(74.4) **76.4**	(61.9) **72.5**	(67.0) **74.9**	(67.8) **68.2**	(58.6) **73.4**
15	(13.1) **40.4**	(15.0) **45.5**	(14.2) **40.2**	(14.6) **53.0**	(14.6) **47.4**	(14.6) **47.7**	(14.3) **45.0**	(15.9) **39.9**	(15.2) **28.6**	(14.6) **43.1**
17	(17.3) **52.0**	(25.1) **52.4**	(23.5) **51.0**	(24.2) **45.2**	(24.6) **45.7**	(24.1) **48.4**	(22.0) **43.7**	(21.1) **44.8**	(19.5) **36.5**	(22.4) **46.6**
23	(22.7) **52.0**	(22.9) **54.2**	(19.9) **53.6**	(20.8) **49.4**	(19.6) **53.9**	(16.6) **54.2**	(20.1) **49.4**	(20.1) **44.4**	(20.8) **34.9**	(20.4) **49.6**
Av.	(36.0) **64.1**	(41.8) **65.7**	(23.2) **63.4**	(46.5) **65.8**	(48.2) **68.6**	(46.8) **66.8**	(40.8) **64.3**	(41.8) **62.7**	(40.4) **57.9**	(40.6) **64.4**
AU	ACC									
	v1	v2	v3	v4	v5	v6	v7	v8	v9	**Global**
1	(25.2) **91.2**	(35.9) **92.9**	(32.9) **92.2**	(56.3) **94.6**	(58.0) **93.4**	(50.0) **95.2**	(72.9) **92.3**	(90.2) **93.7**	(91.5) **94**	(57.0) **93.3**
4	(30.0) **92.1**	(37.2) **94.4**	(14.3) **90.0**	(45.6) **93.2**	(55.4) **97.3**	(62.7) **94.0**	(50.1) **92.6**	(81.0) **92.7**	(**92.0**) 88.1	(52.0) **92.7**
6	(73.2) **81.4**	(72.7) **81.6**	(69.6) **84.5**	(81.3) **84.5**	(82.9) **84.1**	(79.1) **85.9**	(66.7) **84.2**	(47.3) **80.1**	(35.7) **82.7**	(67.6) **83.2**
7	(60.6) **75.3**	(64.2) **75.9**	(53.9) **73.9**	(69.5) **76.3**	(73.8) **79.0**	(72.8) **77.6**	(61.0) **73.8**	(60.4) **74.9**	(61.7) **72.4**	(64.2) **75.5**
10	(56.3) **80.1**	(65.0) **83.1**	(71.1) **81.4**	(74.9) **81.9**	(72.2) **83.4**	(67.5) **83.6**	(61.2) **83.5**	(56.0) **82.8**	(50.3) **80.6**	(63.8) **82.3**
12	(62.7) **87.1**	(70.7) **87.8**	(60.9) **86.0**	(81.0) **87.5**	(81.1) **86.2**	(76.4) **86.2**	(63.6) **83.9**	(53.9) **85.3**	(43.6) **83.4**	(66.0) **85.9**
14	(61.8) **62.5**	(62.6) **64.6**	(51.6) **62.4**	(71.2) **65.0**	(73.0) **71.1**	(**73.0**) 72.7	(60.8) **63.4**	(53.3) **66.6**	(52.4) **54.9**	(62.2) **64.8**
15	(39.8) **86.8**	(21.8) **90.4**	(10.2) **84.7**	(67.8) **92.4**	(21.5) **91.0**	(12.1) **91.4**	(36.8) **90.6**	(28.5) **89.8**	(37.8) **75.9**	(30.7) **88.1**
17	(69.9) **83.4**	(46.8) **86.5**	(24.6) **85.0**	(74.1) **76.1**	(59.8) **85.8**	(52.2) **84.9**	(46.4) **79.8**	(30.6) **82.0**	(32.3) **76.4**	(48.5) **82.2**
23	(27.1) **87.5**	(27.5) **88.5**	(55.4) **87.3**	(50.8) **87.3**	(46.6) **87.8**	(59.4) **87.9**	(34.9) **86.4**	(18.9) **82.3**	(15.4) **76.7**	(37.3) **85.7**
Av.	(50.7) **82.7**	(50.5) **84.6**	(44.4) **82.7**	(67.3) **83.9**	(62.4) **85.9**	(60.5) **85.9**	(55.4) **83.1**	(52.0) **83.0**	(51.3) **78.5**	(54.9) **83.4**

Table 4: **Results over FERA17 Validation, using F1-score and Accuracy (ACC)**. Comparison of our approach against the state-of-the-art method [43] (presented in parentheses).

Evaluation on Test Set

We include the results obtained during the FERA17 Challenge, where the test server was made available from February 8 to March 1 2017 with a limit of 5 submissions per group. The challenge rules required participants to upload their source code to the FERA17 server, and the organizers ran the models on an undisclosed test set. Following the challenge setup, we report F1-score and ACC in Table 4 and 5, for Val and Test sets. Our final submitted system is represented in Fig. 2. As can be observed, our model sets a new benchmark in the problem of multi-view action unit detection by improving with an **14%** in F1-score over the baseline approach in the Test set.

Efficiency

We measure the average execution time for the final multi-view system as follows: For a video with 500 frames, it takes on average 12 seconds to transform it into frames. The Optical Flow takes about 2 minutes producing the whole set of OF frames. Once the RGB and OF frames are complete, the view detector takes around 2 seconds predicting the corresponding view of the video from a frame sub-sampling.

AU	F1-score (Baseline [43]) Ours									
	v1	v2	v3	v4	v5	v6	v7	v8	v9	**Global**
1	(11.4) **33.7**	(14.7) **30.5**	(15.4) **28.8**	(16.0) **35.0**	(18.9) **32.3**	(17.3) **28.1**	(11.8) **34.0**	(3.50) **27.5**	(0.50) **28.7**	(12.2) **30.9**
4	(4.40) **22.6**	(4.30) **29.2**	(5.50) **13.2**	(4.10) **18.9**	(3.20) **14.9**	(**3.70**) 2.90	(4.70) **14.7**	(4.10) **19.3**	(1.50) **19.0**	(3.90) **16.6**
6	(55.0) **79.8**	(60.2) **80.3**	(20.4) **80.4**	(70.4) **79.8**	(**80.8**) 79.2	(78.7) **80.1**	(56.5) **79.7**	(64.3) **80.4**	(61.5) **79.7**	(60.9) **79.9**
7	(74.2) **82.0**	(77.0) **85.2**	(32.9) **85.0**	(78.8) **83.3**	(**83.7**) **83.7**	(**84.0**) 83.5	(75.8) **83.2**	(77.7) **82.8**	(77.9) **83.3**	(73.6) **83.6**
10	(66.0) **80.7**	(77.2) **83.8**	(67.0) **81.4**	(83.9) **85.9**	(83.1) **87.8**	(81.2) **85.8**	(66.3) **86.1**	(76.7) **84.7**	(75.8) **84.3**	(75.2) **84.6**
12	(60.9) **82.3**	(70.2) **81.8**	(26.1) **83.7**	(81.1) **84.0**	(80.0) **84.5**	(79.8) **84.1**	(60.9) **79.3**	(69.8) **82.1**	(67.8) **81.9**	(66.3) **82.7**
14	(59.3) **77.4**	(61.1) **78.9**	(10.9) **78.0**	(66.3) **78.1**	(77.6) **79.1**	(**80.3**) **80.3**	(63.1) **76.8**	(74.5) **77.8**	(74.9) **77.3**	(63.1) **78.2**
15	(16.2) **23.2**	(**23.2**) 21.1	(22.7) **38.8**	(14.7) **26.6**	(22.5) **28.0**	(22.9) **24.1**	(23.0) **26.4**	(23.6) **30.2**	(23.7) **29.1**	(21.4) **28.0**
17	(16.4) **50.1**	(34.6) **48.4**	(30.6) **52.5**	(17.7) **43.2**	(30.6) **43.5**	(31.1) **48.7**	(22.9) **42.3**	(26.0) **37.9**	(26.6) **37.8**	(26.3) **44.3**
23	(37.3) **46.5**	(37.1) **51.4**	(31.3) **51.7**	(27.8) **50.6**	(31.7) **54.3**	(26.8) **51.3**	(31.6) **42.4**	(36.6) **44.3**	(36.8) **44.7**	(33.0) **48.6**
Av.	(40.1) **57.8**	(46.0) **59.1**	(26.3) **59.4**	(46.1) **58.5**	(51.2) **58.7**	(50.6) **56.9**	(41.6) **56.5**	(45.7) **56.7**	(44.7) **56.6**	(43.6) **57.7**

AU	ACC									
	v1	v2	v3	v4	v5	v6	v7	v8	v9	**Global**
1	(25.1) **92.4**	(33.1) **91.3**	(40.3) **90.6**	(36.1) **90.3**	(50.5) **87.7**	(43.8) **89.9**	(64.1) **89.2**	(**91.2**) 88.0	(**93.2**) 88.7	(53.0) **89.8**
4	(39.9) **97.3**	(31.4) **97.6**	(18.5) **93.8**	(49.9) **94.7**	(66.2) **87.9**	(69.8) **94.9**	(46.0) **92.7**	(85.7) **91.5**	(**93.6**) 91.2	(55.7) **93.5**
6	(63.7) **81.5**	(70.0) **81.9**	(60.1) **81.5**	(77.9) **80.1**	(**83.2**) 79.0	(80.5) **81.6**	(62.1) **80.9**	(52.6) **81.5**	(45.6) **80.8**	(66.2) **81.0**
7	(61.4) **76.4**	(66.2) **80.5**	(45.9) **80.3**	(74.6) **77.1**	(**78.6**) 78.4	(**79.0**) 78.5	(64.3) **75.9**	(64.0) **75.8**	(63.9) **76.6**	(66.4) **77.7**
10	(59.5) **76.8**	(67.8) **81.5**	(65.6) **79.2**	(78.6) **81.9**	(75.4) 84.3	(71.8) **81.9**	(58.7) **82.4**	(64.3) **81.4**	(61.9) **80.9**	(67.1) **81.1**
12	(60.5) **81.8**	(70.0) **82.6**	(54.8) **83.6**	(79.5) **82.7**	(76.0) **83.5**	(75.6) **83.1**	(59.8) **79.0**	(57.7) **82.5**	(52.3) **82.4**	(65.1) **82.4**
14	(58.3) **66.5**	(60.7) **70.4**	(42.3) **68.7**	(65.2) **68.7**	(**73.6**) 72.6	(**75.7**) 73.2	(58.1) **67.7**	(59.9) **68.2**	(59.9) **67.1**	(61.5) **69.2**
15	(45.0) **82.4**	(21.3) **86.1**	(13.9) **84.2**	(62.0) **84.6**	(25.5) **86.7**	(17.0) **85.1**	(35.9) **82.9**	(24.5) **81.2**	(33.6) **79.7**	(31.0) **83.6**
17	(71.5) **83.0**	(56.4) **84.3**	(35.3) **85.1**	(75.3) **75.4**	(66.5) **82.5**	(55.9) **82.3**	(47.6) **74.0**	(24.9) **73.1**	(36.1) **72.1**	(52.2) **79.0**
23	(32.5) **72.2**	(27.5) **77.0**	(53.3) **75.9**	(58.1) **75.5**	(57.4) **79.1**	(67.7) **72.6**	(39.6) **72.9**	(28.3) **71.4**	(24.2) **72.1**	(43.2) **74.3**
Av.	(51.7) **81.0**	(50.4) **83.3**	(43.0) **82.3**	(65.7) **81.1**	(65.3) **82.2**	(63.7) **82.3**	(53.6) **79.8**	(55.3) **79.5**	(56.4) **79.2**	(56.1) **81.2**

Table 5: **Results over FERA17 Test set, using F1-score and Accuracy** . Comparison of our approach against the state-of-the-art method [43] (presented in parentheses).

For the final step, it takes roughly 8 more minutes to forward every frame of the video through the 10 AUs detectors. In summary, it takes less than 1 second to process each frame once the OF is computed.

4.3. Qualitative Results

We present qualitative results over BP4D dataset using the visualization techniques of Zeiler and Fergus [49], and Yosinski [48].

One of Zeiler's insights is to occlude parts of the input image in order to look for the classification probability this region is supplying to the full image. Fig. 4 and 5 show this approach over several images for different AUs. We can observe that our approach emphasizes very specific regions of the face (third row) using thus local information for each action unit despite analyzing the face hollistically.

Additionally, Yosinski's method creates synthetic images as input in order to maximize the output of one specific neuron; for instance, we maximize the output of the binary one. Similar to Zeiler's, Fig. 6 shows that AU1 and AU2 models are focused on the eyes, thus AU12 and AU15 representations are looking for patterns over the mouth region, confirming thus the specificity of these models.

5. Conclusions

We have presented HydraNets and AUNets for recognizing 12 different facial action units. Our method takes advantage of the power of CNNs for large-scale classification problems, and is capable of detecting multiple action units simultaneously. Our method uses the temporal information OF to enhance performance. There is a trade-off between efficiency and performance with HydraNets and AUNets, yet both approaches present favorable results when compared against state-of-the-art methods, and our final multi-view system provides a large improvement in performance in a challenging benchmark. At the core of our approach lies a flexible and modular architecture that can easily incorporate new action units. In order to promote further research in multi-view action unit detection, we will make our source code, models, and results publicly available.

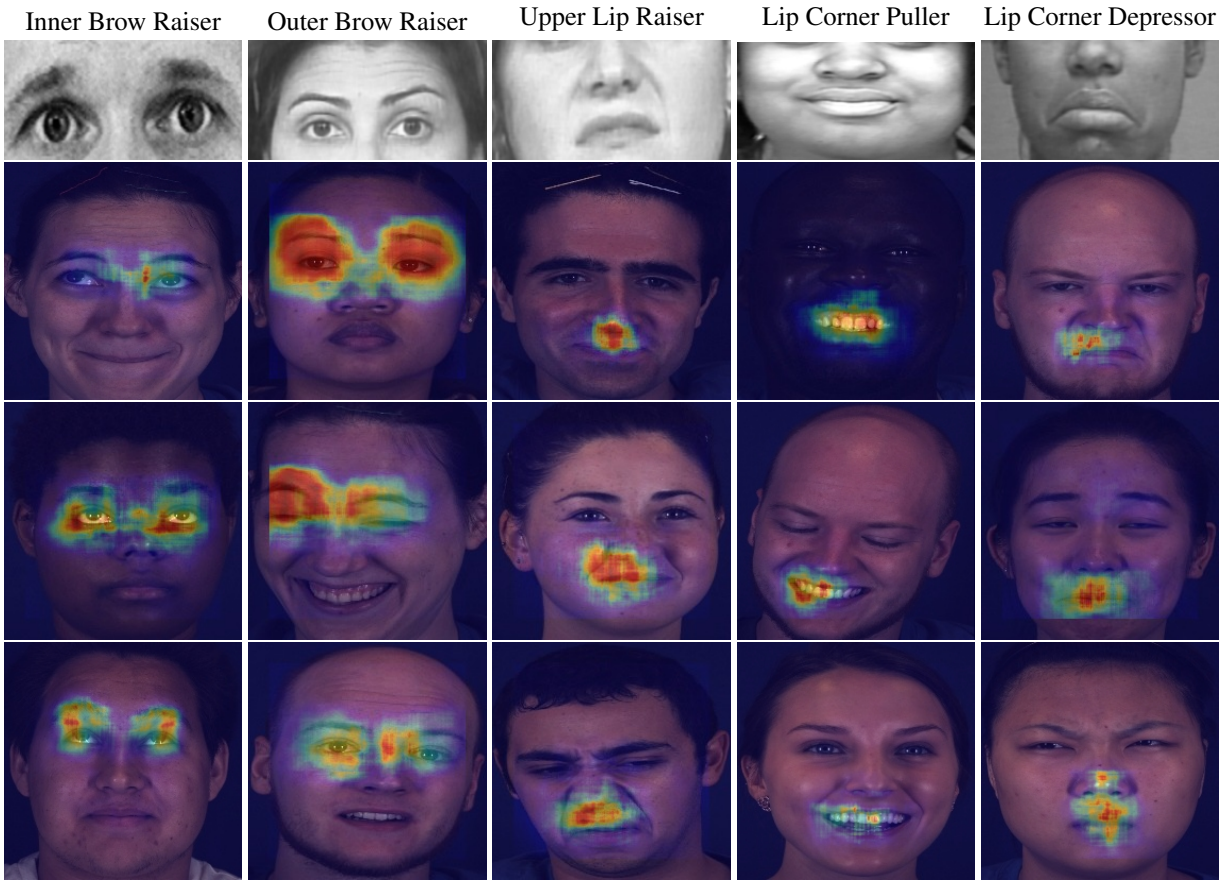

Figure 4: Zeiler's method [49] for network visualization. The top row presents 5 different Action Units [7]. The heat maps highlight the most important regions in the human face for each specific Action Unit (blue: less important, red: more important).

References

[1] T. Almaev, B. Martinez, and M. Valstar. Learning to transfer: Transferring latent task structures and its application to person-specific facial action unit detection. In *ICCV*, 2015. 2, 3

[2] T. Brox and J. Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *TPAMI*, 2011. 6

[3] F. Caba Heilbron, V. Escorcia, B. Ghanem, and J. Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 3

[4] R. Chaudhry, A. Ravichandran, G. Hager, and R. Vidal. Histograms of oriented optical flow and binet-cauchy kernels on nonlinear dynamical systems for the recognition of human actions. In *CVPR*, 2009. 3

[5] G. Chéron, I. Laptev, and C. Schmid. P-cnn: Pose-based cnn features for action recognition. In *ICCV*, 2015. 2

[6] W.-S. Chu, F. De la Torre, and J. F. Cohn. Modeling spatial and temporal cues for multi-label facial action unit detection. *arXiv preprint arXiv:1608.00911*, 2016. 3, 5, 6, 7

[7] CMU. Facial action coding system, 2002. 10, 11

[8] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *ECCV*. 1998. 3

[9] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. In *ICCV*, 2016. 3

[10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 2, 3

[11] P. Ekman and W. V. Friesen. Facial action coding system. 1977. 1

[12] P. Ekman and E. L. Rosenberg. *What the face reveals: Basic and applied studies of spontaneous expression using the Facial Action Coding System (FACS)*. Oxford University Press, USA, 1997. 1

[13] S. Eleftheriadis, O. Rudovic, and M. Pantic. Multi-conditional latent variable model for joint facial action unit detection. In *ICCV*, 2015. 2, 3

[14] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 3

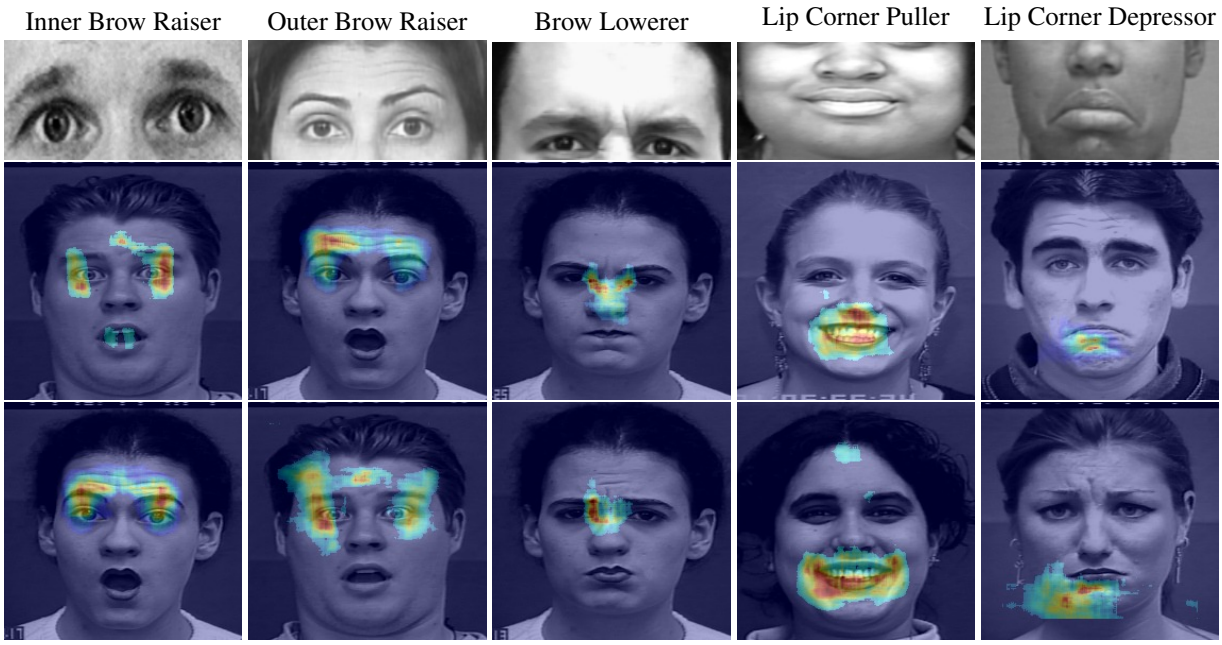

Figure 5: Zeiler's method [49] for network visualization. The top row presents 5 different Action Units [7]. The heat maps highlight the most important regions in the human face for each specific Action Unit (blue: less important, red: more important).

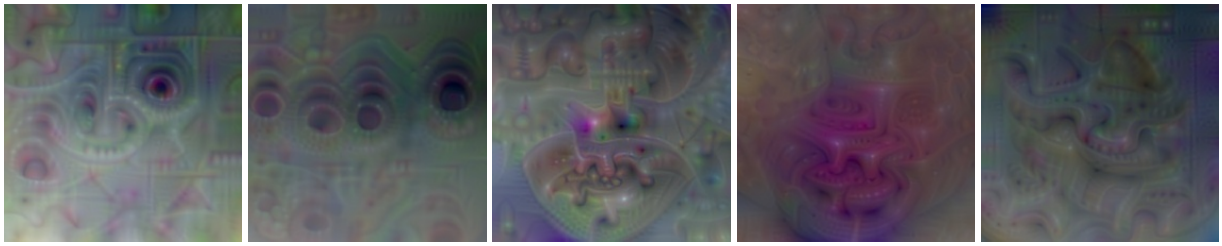

Figure 6: Yosinski's method [48] for network visualization. This approach shows synthetic images that maximize the output of the network for AU 1,2,10,12,15.

[15] G. Gkioxari and J. Malik. Finding action tubes. In *CVPR*, 2015. 4, 6

[16] A. Gudi, H. E. Tasli, T. M. den Uyl, and A. Maroulis. Deep learning based facs action unit occurrence and intensity estimation. In *11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition*, 2015. 2, 3

[17] S. Happy and A. Routray. Automatic facial expression recognition using features of salient facial patches. *IEEE transactions on Affective Computing*, 6(1):1–12, 2015. 1

[18] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 3

[19] B. K. Horn and B. G. Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 3

[20] Itseez. Open source computer vision library. https://github.com/itseez/opencv, 2015. 6

[21] S. Jaiswal and M. Valstar. Deep learning the dynamic appearance and shape of facial action units. In *WACV*, 2016. 2, 3

[22] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM*, 2014. 7

[23] V. G. Kaburlasos, S. E. Papadakis, and G. A. Papakostas. Lattice computing extension of the fam neural classifier for human facial expression recognition. *IEEE Transactions on Neural Networks and Learning Systems*, 24(10):1526–1538, Oct 2013. 5

[24] S. E. Kahou, P. Froumenty, and C. Pal. Facial expression analysis based on high dimensional binary features. In *ECCV*, 2014. 1

[25] S. Kaltwang, S. Todorovic, and M. Pantic. Latent trees for estimating intensity of facial action units. In *CVPR*, 2015. 2,

3

[26] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 2, 3, 5

[27] W. Li, F. Abtahi, Z. Zhu, and L. Yin. Eac-net: A region-based deep enhancing and cropping approach for facial action unit detection. *arXiv preprint arXiv:1702.02925*, 2017. 2, 3, 5, 6, 7

[28] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2

[29] M. Liu, S. Li, S. Shan, and X. Chen. Au-aware deep networks for facial expression recognition. In *Automatic Face and Gesture Recognition (FG)*, 2013. 3

[30] M. Liu, S. Li, S. Shan, R. Wang, and X. Chen. Deeply learning deformable facial action parts model for dynamic expression analysis. In *ACCV*, 2014. 3

[31] P. Liu, S. Han, Z. Meng, and Y. Tong. Facial expression recognition via a boosted deep belief network. In *CVPR*, 2014. 3

[32] Y.-J. Liu, J.-K. Zhang, W.-J. Yan, S.-J. Wang, G. Zhao, and X. Fu. A main directional mean optical flow feature for spontaneous micro-expression recognition. *IEEE Transactions on Affective Computing*, 7(4):299–310, 2016. 2, 3

[33] D. G. Lowe. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, volume 2, pages 1150–1157. Ieee, 1999. 3

[34] P. Lucey, J. F. Cohn, T. Kanade, J. Saragih, Z. Ambadar, and I. Matthews. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *CVPR*, 2010. 1, 5

[35] M. Mavadati, P. Sanger, and M. H. Mahoor. Extended disfa dataset: Investigating posed and spontaneous facial expressions. In *CVPR Workshops*, 2016. 3

[36] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *ICCV*, 2015. 3

[37] T. Ojala, M. Pietikainen, and D. Harwood. Performance evaluation of texture measures with classification based on kullback discrimination of distributions. In *Pattern Recognition, 1994. Vol. 1-Conference A: Computer Vision & Image Processing., Proceedings of the 12th IAPR International Conference on*, volume 1, pages 582–585. IEEE, 1994. 3

[38] X. Peng and C. Schmid. Multi-region two-stream r-cnn for action detection. In *ECCV*, 2016. 4, 6

[39] A. Savran, N. Alyüz, H. Dibeklioğlu, O. Çeliktutan, B. Gökberk, B. Sankur, and L. Akarun. Bosphorus database for 3d face analysis. In *Biometrics and Identity Management*, pages 47–56. Springer, 2008. 5

[40] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2, 3, 5

[41] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 2, 3, 5

[42] M. F. Valstar, M. Mehu, B. Jiang, M. Pantic, and K. Scherer. Meta-analysis of the first facial expression recognition challenge. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 42(4):966–979, 2012. 5

[43] M. F. Valstar, E. Sánchez-Lozano, J. F. Cohn, L. A. Jeni, J. M. Girard, Z. Zhang, L. Yin, and M. Pantic. Fera 2017-addressing head pose in the third facial expression recognition and analysis challenge. *arXiv preprint arXiv:1702.04174*, 2017. 1, 3, 6, 7, 8, 9

[44] Z. Wang, Y. Li, S. Wang, and Q. Ji. Capturing global semantic relationships for facial action unit recognition. In *ICCV*, 2013. 3

[45] S. "Xie and Z. Tu. Holistically-nested edge detection. In *ICCV*, 2015. 3

[46] J. Yang, B. Price, S. Cohen, H. Lee, and M.-H. Yang. Object contour detection with a fully convolutional encoder-decoder network. In *CVPR*, 2016. 3

[47] L. Yang, P. Luo, C. Change Loy, and X. Tang. A large-scale car dataset for fine-grained categorization and verification. In *CVPR*, 2015. 2

[48] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson. Understanding neural networks through deep visualization. In *ICML*, 2015. 9, 11

[49] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014. 9, 10, 11

[50] J. Zeng, W.-S. Chu, F. De la Torre, J. F. Cohn, and Z. Xiong. Confidence preserving machine for facial action unit detection. In *ICCV*, 2015. 2, 3

[51] X. Zhang, L. Yin, J. F. Cohn, S. Canavan, M. Reale, A. Horowitz, P. Liu, and J. M. Girard. Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database. *Image and Vision Computing*, 32(10):692–706, 2014. 1, 5, 6

[52] Z. Zhang, J. M. Girard, Y. Wu, X. Zhang, P. Liu, U. Ciftci, S. Canavan, M. Reale, A. Horowitz, H. Yang, et al. Multimodal spontaneous emotion corpus for human behavior analysis. In *CVPR*, 2016. 3, 6

[53] K. Zhao, W.-S. Chu, F. De la Torre, J. F. Cohn, and H. Zhang. Joint patch and multi-label learning for facial action unit detection. In *CVPR*, 2015. 2, 3, 5, 6, 7

[54] K. Zhao, W.-S. Chu, and H. Zhang. Deep region and multi-label learning for facial action unit detection. In *CVPR*, 2016. 2, 3, 5, 6, 7